

More Than a Show: Using Personalized Immersive Theater to Educate and Engage the Public in Technology Ethics

Michael Skirpan
University of Colorado
Boulder, CO
michael.skirpan@colorado.edu

Jacqueline Cameron
University of Colorado
Boulder, CO
jacqueline.cameron@colorado.edu

Tom Yeh
University of Colorado
Boulder, CO
tom.yeh@colorado.edu

ABSTRACT

Devising strategies to engage the public in discussions around the design and development of technology is critical to building a future that works for everyone. This paper presents a novel case study, an immersive theater experience, "Quantified Self," that combines aspects of design fiction and user enactments to construct a public engagement opportunity about technology ethics. Our audience supplied their social data (Facebook, Twitter...) and received a personalized experience where they interacted with a narrative and technology exhibits. We used a design model targeting goals of engagement, education, and discussion. Here we overview the design and production of Quantified Self and report on the results (240 participants over 6 performances) and findings from audience surveys (n=179/240) and cast/crew interviews (n=15/22). We found our approach attracted a wide audience interested in different elements of the show. Affordances and challenges of our model are discussed in detail.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords

Design Methods; Immersive Theater; Speculative Design; Ethics; Public Engagement; Research through Design; Design Fiction; Enactments; Privacy

INTRODUCTION

Choices being made about how technology is designed, developed, and used shape the character of public and private life. Yet, there are still many open questions about the purpose of technology in human life and society: What should and should not be done with our personal data? [1, 45, 14] How do we apply laws and protect rights in the space of technology? [26, 27] To what extent should we allow choices by technologists to disrupt social norms? [19, 1, 22] In this techno-society, it is imperative that we broaden participation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-5620-6/18/04... 15.00

DOI: <https://doi.org/10.1145/3173574.3174038>

in the discussions around technology's development to support a negotiation toward a future that works for everyone; not just a small elite group of technologists and policy makers.

Creating an opportunity for this broad discussion is truly difficult. Not only does it demand engaging diverse groups to participate in discussion and reflection, but also one must communicate complex issues in a format accessible to non-experts. To date, perhaps the most successful format for wide engagement in technological and scientific problems is to represent them in artistic and diegetic manners [24, 39, 25, 38, 5, 34, 35]—e.g., sci-fi, theater, and concept art. Researchers too use design fiction [7, 46, 48], design ethnography [30], user enactments [40, 20], and other forms of speculative art [18] to pose questions about the future. It is within this trend of combining art and technology to raise opportunities for discourse and research that inspired our project, "Quantified Self: Immersive Theater and Data Experience." In this paper, we present our approach, experience, and findings as a case study toward how to do futures research, particularly in the space of technology ethics, at scale.

Our goal was to design a public engagement program that raised awareness and discussion around what it means for mass amounts of personal data to be owned and used by third parties. The motivation was to go beyond questions of privacy and focus on what it may be like to live in a world where all our data is overtly used. Will people like that reality? Will it be fair? What control should we have over our own information?

Inspired by recent HCI work on enactments and design fiction and the gripping power of contemporary immersive theater [12], we embarked on a year-long production that would engage, educate, and provoke dialogue about the future of data. The result was Quantified Self which drew over 240 audience members of diverse backgrounds to six performances and was co-produced by a cross-disciplinary team of 22 students representing computer science, electrical engineering, theater, fine arts, and physics.

RELATED WORK

Why We Need More Conversations about Our Future

It should come as no surprise to HCI researchers that there is a desperate need to educate and converse with the public on the nature and potential futures of technology, in particular the uses of Big Data. Recent events and research raise

red flags that all the new ways of using data may not lead to shared prosperity. The groundbreaking work done at Pro Publica highlighted how machines are capable of being discriminatory much like humans [26]. Further work has made this risk evident as police face databases are biased toward African American faces [28], word embeddings have proven to encode gender stereotypes [10], and even behavioral advertising is codifying unfortunate racial differences [47].

There has further been attention paid to the new powers of data through predictive inference. 2016 in particular made evident the reality of algorithms shaping our society [19] and affecting public discourse [14]. And while prescient members research community may have identified warning signs [50, 15, 11, 41] for problems these changes may create, it's unclear technologists are really connected to the best interests of the public and their users [22].

Past research has shown that the mass move toward personalization is seen as unwanted or even worrying by some users [49, 1]. Even studies done by Google engineers availed that users do not see themselves as receiving the benefit of mass data acquisitions and sharing [6]. Moreover, we should not be surprised by the umbrage taken by certain communities [45] when plenty of HCI research points to rampant misunderstandings around the nature of the terms of service that legally bind us and the underlying technologies that track, categorize, and target us [37, 23, 32].

Approaches to Studying Problems of the Future

While we have come to understand the problems of ignorance and public concern with technology, researchers have continually worked to discover our best ways to communicate and study the problems. There is a long history that suggests sci-fi and other forms of speculative art are critical vehicles for reaching the broader public and raising awareness about the problems the future may bring [18, 24, 51, 8, 3, 44]. This history of interaction between technology and art has given rise to an appreciation of the influence and impact by researchers [35, 36, 29, 38, 31]. Knowing the affordances of art for technological and scientific discourse, Bruce Sterling formalized the idea of design fiction in order to establish creative methods for researchers to employ imagination and pose questions about technology where social, political, and emotional content is integrated [46, 7].

Years later, we see a rich array of approaches being explored to utilize art as a means toward critique [18, 3] and conceptualizing where certain areas are headed [17]. Though not necessarily employing art, future studies is an area that has attempted to formalize the study of possible futures [33] and the potentials for utopia and dystopia [13, 16]. However, within the space of art and design, many incredible projects have emerged that have affordances for research.

The notion of a design or anticipatory ethnography [30] has brought qualitative methods closer to art and artists, assessing details about worlds that do not yet exist and the thought that goes into making them. Most influential on our thinking has been the creative approach of integrating theatrics into the study of how people feel about the future. Enactments

[40, 20] and 'lived informatics' [21] allow us to take scenarios and design prototypes a step further to construct sites of research [53]. That is, by placing people into environments and scenes representing possible futures, we can leverage the role of improvisation to see how people act and feel within these settings.

It is in the overlap of these areas—design fiction and user enactments—we found potential for a project that can at once engage and educate while also allowing for discussion and listening. Thus, it was our goal to meld together some of these pioneering methods to push the potential for broad, inter-community engagement and research regarding the future of data use and ownership. Our unique contribution was trying to create something that could be scaled to large audiences of diverse backgrounds. Even if we know fiction and enactments have the potential for reflection and inquiry by participants, it is not easy or obvious how these efforts could be formatted for broad accessibility. That is, we consider the balance between engagement and entertainment, detailed technical content and artistic simplification as well as audience interaction and narrative structure. With Quantified Self as our case study, we hope to enrich the literature on what might be the best way to engage large audiences of diverse backgrounds.

DESIGN SPACE

We begin by offering a heuristic for thinking about the design space related to the problems discussed above: informing the public, engaging a wide audience, provoking dialogue, and learning from our participants. The heuristic was designed to simplify thinking about other tech/art projects that were already achieving educational goals while building opportunity for other kinds of discovery such as research, inter-group conversation, and heightened individual awareness of technology's impacts.

The dimensions we sought to map were: 1) The degree to which art is being used to represent human-centric problems such as emotions, relationships, and qualifying experience. 2) The degree to which technology is being used to represent quantitative content such as mined data, live sensing, and mathematical simulation. 3) The degree to which the overall content is fixed prior to engaging an audience. 4) The degree to which content is improvised and adapted, allowing audiences to inject agency into the final experience.

Our goal was to consider how our program could leverage a highly balanced approach, designing a piece that involved real technology and artistic abstraction, fixed structure and improvisation. We hoped such a balance would lead to a type of performance that is engaging, sustainable (i.e., not needing a ton of expert attention for others to run it), and scalable (i.e., could be modified, repeated, and take in larger audiences). Let's consider this in comparison to a sci-fi think-piece such as the movie "Her." The movie references recognizable technologies such as AI, voice recognition, and portable screens; though, it was also abstracted into an artistic view where human-computer relationships were very advanced along with our cities and personal technologies. "Her," is fantastic for considering love in light of technology, but as a fixed narrative, any reflection, discussion, or

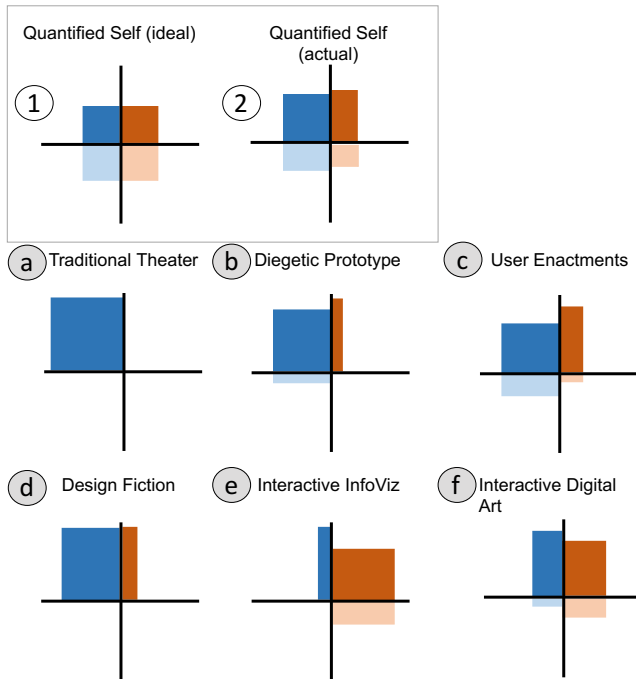


Figure 1. The goal of Quantified Self is to afford a balanced experience across all dimensions of our heuristic (1). We came close to meeting this goal (2). Below are examples of other projects using our heuristic for considering the affordances of different designs (a-f).

agency was left for viewers after the show. We wanted our audience to be able to feel part of the story and make real decisions with real technology that had consequence.

More in line with our vision, a project by David Chatting, *Runner Spotters* [20], was motivating. He got actors interacting with designed technology artifacts in partially-improvised scenes to allow a designer to speculate on what kind of social constructs may emerge from a new technology. The result was akin to a performative design fiction since there was little real technology used and the scope of interaction was fixed to a very small group of actors whose interactions were used to make the final short film and design fiction. While we loved the possibilities presented by the work, it relied heavily on a fixed structure and the art/aesthetics of technology over real technology.

Chris Elden’s work on *Metadating* [20] provided another useful example for us. *Metadating* was a one-off speeddating event where the audience was encouraged to fill out data profiles to use as their information exchange for real social interaction. The ability for the audience to make what they wanted of the event and connect their data into something personal and consequential was very influential on us. However, it was a highly catered event with very little fixed structure to ensure certain ideas were raised and in the end only served 11 individuals. A novel addition we wanted to an event like this was for a stronger story that could focus audience reflection on themes and ideas while allowing more people to participate more easily.

Finally, Odom’s work on user enactments [40] was the closest to our vision. William Odom has used the framework of the enactment to allow pre-structured scenarios involving possible future technologies to be acted out by participants where crucial moments are left improvised. We were intrigued by his approach to putting the user in the driver seat with enough information to get going, but open endings to offer real agency and a site for interesting futures research around hard topics like social norms. Where we wanted to iterate on Odom’s work was by opening up participation beyond the constraint of a set of scripted characters, but open up a world where few or many people could interact at once. We also hoped to make a more seamless way for an audience member to get a personalized experience in these scenarios without having to put time into collecting data from each group individually.

In contrast these pieces, our performance attempted a set of novel additions. 1) We wanted our event to have a structured story that was easy for communication and focused on issues, such as is found in a movie or traditional theater, while having enough improvisational possibility for discovery. 2) Our goal was to allow actors to have flexible interactions with technology, like *Runner Spotters*, but further open up a world where audience-actor-technology interactions were all possible. 3) The project was meant to be able to adapt to larger audiences than recent work in enactments has achieved. This meant we needed an automated approach to data collection and a large enough world that many interactions, improvised and planned, could be happening at once. As Figure 1 shows, we hoped these factors would provide a balance between the amount of fixed vs. improvised content as well as between real technology vs. artistic abstraction.

APPROACH AND METHODS

To obtain a balanced experience as laid out above, we developed “Quantified Self”: an immersive theater piece where audience members were part of the performance. They were able to walk freely through the set, converse with actors and other audience members, and interact with the set and story using their own volition and at their own pace. Technology exhibits, or “companions” in our story line, were installed throughout the set, offering personalized interactions using social media data that our audience shared at ticketing.

High-Level Design

Beyond the artistic presentation, to further increase engagement, we aimed to maximize interactivity and personalization. Much like our contemporary technologies, we wanted a theater performance that adapted to the audience. Designing an open format also promoted discussion by creating communication pathways between the different communities coming together. Audience members could talk to friends, to strangers (often from a different background), actors, and, during the talk back, to our production team. Different perspectives in these conversations were embedded into the story and could happen organically by attracting a broad audience.

In terms of content, the show incorporated four primary educative themes on technology and data ethics:

Approximation of Self (A) Our goal with this issue was to discuss the controversial idea that behavioral metrics and online presentations are enough to capture the nature of a person. Of course, in actuality, human attributes quantified and aggregated online are approximations using statistical methodologies. However, the results of these methods have serious consequences on the shape of online and offline experiences [42]. In what cases is it ethical for them to be used to make important decisions about a person or to schedule and shape their activities?

Data Ownership and Privacy (O) Data is primarily owned by the companies providing the service, leaving the user with little control over how their data is used. Where law could be a safeguard, we often find that people are blindly clicking and agreeing to contractual agreements they largely do not understand [4, 9, 23]. When browsing websites, tracking can occur without opt-in consent. Companies offer free services by being able to buy and sell these everyday interactions. What should companies have the right to do with our data? What transparency should there be regarding these uses?

Presentation of Data (R) This issue looked at how people may be influenced or manipulated by the presentation of information that seems plausible but may be fallacious [52]. How does the presentation of data impact how you interpret it? Can showing quantified results breed undeserved trust?

Personalization (P) The issue of personalization wrestles with the current trend in technology that values interfaces and systems that use current and past information to tailor an experience to a particular user. Critically, there are times where personalization may create privacy concerns by exposing or implying information about someone or be an undesired hindrance such as when experiencing a search bubble phenomenon [49, 19]. How much of our world should be personalized? Would we like a highly personalized world?

Design Process

Quantified Self was developed through an iterative design process over a one-year period (see Figure 2). The production lead developed the initial idea and storyline, setting the theme and goals. The co-producers worked with the lead to develop the sign-up process, technology back-end, and the initial script. Three months prior to the show, the broader team was finalized (Figure 2.1).

To promote diverse discussions within the show, we had a team that crossed a breadth of technological and artistic backgrounds to support in the development and implementation. Our production team primarily consisted of 22 university students, crossing 7 different departments, with additional support from industry partners and faculty advisors (Figure 2.1). The production was led by a third year PhD student in computer science and co-produced with two other PhDs (one studying theater; the other interdisciplinary technology) and an industry data scientist. The leads were supported by 19

other students (15 undergraduate, 4 graduate-level) divided into two main teams: technical and theatrical.

Undergraduate students were brought on the technical team to design and develop exhibits. The complete theatrical team started rehearsals and helped develop the characters through practiced improv. Overall, there were 5 rehearsal days (Figure 2.3) where the technical team worked directly with the theatrical team, sharing technical knowledge with the actors and receiving feedback on their exhibit designs.

Building The Experience

Script

The first piece of the show developed was our script, which was written by the project lead. Here we look at the two primary considerations that went into writing the script: the perspectives it embedded and how it supported the desired experience design.

The show had an overarching narrative following an ethical conflict within a famous tech company, DesignCraft. Immediately upon signing up for the show, participants were invited to a party for their supposed friend, Amelia, who was a star employee at DesignCraft. As the story unravels, they learn that Amelia is an experimental AI created using their personal data, who, herself, has begun grappling with the ethics of how the company uses her and its vast trove of data.

Within this broader plot arc, main characters were written to offer contrasting perspectives on our issues. Don, the CEO of DesignCraft, represented a business and innovation perspective. Lily, the chief data scientist of DesignCraft, held scientific and humanitarian views on the possibilities of Big Data while struggling with some privacy concerns. Felicia, an ex-DesignCraft employee, offered a critical lens of technology infiltrating and destroying the best parts of human relations. Evan, a hacker, saw technology as an opportunity for exploitation and intended to similarly use it to exploit DesignCraft. Amelia, a humanoid AI, struggled with the idea of being merely an instrument for technology and the artificiality of knowing people only through data. Felicity, an FBI agent, believed data could support a more secure society. Bo, the chief marketing officer at DesignCraft, felt strongly that technology was entertaining, useful, and enjoyable and was willing to make this trade-off for any privacy concern. Finally, Veronica, a reporter, was concerned about the politics and intentions of the companies working with everyone's personal data.

Throughout the experience, participants were free to observe and follow different elements of the story, along with potentially discovering that they were actually involved in certain parts of the story. Beyond observing, the script called for actors to engage the audience members in open dialogue through a list of questions and unscripted-yet-topical scenes. During the rehearsal period, actors received training in data issues to be able to have informed discussions with the audience.

Following our heuristic to balance improvisation with fixed structure, we divided the show into four acts, two of which had a unified performance and two of which had multiple

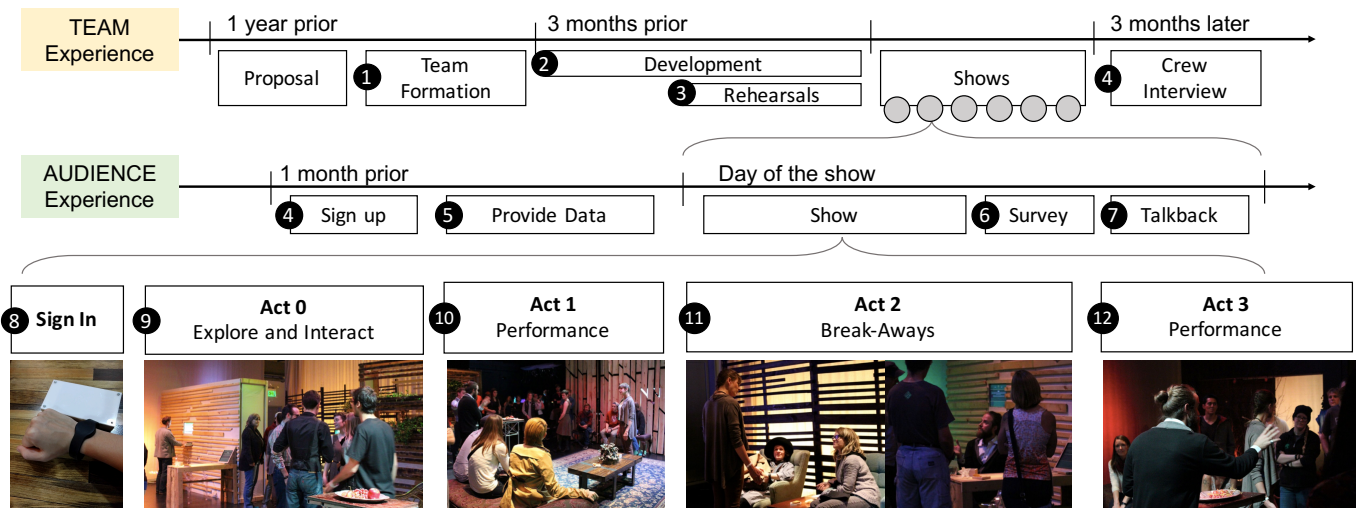


Figure 2. Time line for Quantified Self production.

scenes happening in different locations. In Act 0 (Figure 2.9), after signing in, audience members entered Amelia’s apartment (the stage), where they could freely use the technology exhibits and talk with actors (who may or may not be revealed as such) and each other. This was followed by Act 1 (Figure 2.10), a scripted scene where the main characters set the narrative conflict of Amelia realizing she was an AI-based upon audience data. Act 2 (Figure 2.11), was a series of break-aways, where audience members were encouraged to follow a character to hear a different side of the story. The story culminated in Act 3 (Figure 2.12), with a final scripted dialogue.

A difficult factor in writing the script was deciding how to fit technology fit the world. Since the script was written prior to the creation of the technical artifacts, we hoped to leave open possibilities for iteration and not require technical feats that might have been unfeasible given the timeline and budget. In the end, two (of the final ten) technological artifacts were explicitly built into the script allowing the remaining artifacts to be malleable.

Exhibits

As participants signed up for a ticket (Figure 2.4), they had the option to share their personal data with the show (Figure 2.5). Upon arrival, they were checked in to the party and given a Biobracelet (an RFID wristband) (Figure 2.8) that allowed users to sign into each of the 11 digital “companions” within the set. “Companions” were the name given to the fictitious product line DesignCraft created and provided the plot basis for the interactive exhibits. A digital companion is an interactive set piece which uses participants online or in-person data (e.g., Facebook, Google, Twitter) to create personalized games and experiences, aimed at demonstrating many possible uses of personal data from mimicking online dating to the use of data in a job interview.

We varied the design of our 12 companions along a series of factors to provide participants with a range of experiences and possible levels of interaction. Factors included number of

players, public vs. private play, educative theme (see High-Level Design), type/amount of personal data, and digital vs. non-digital. Two exemplars are Own Up and Discover Your Inner Desire (see Table 1 for full list).

Own Up was structured as a collaborative game for 2-4 players sitting around a screen, embedded face up in a table, with a button on each side. After all players checked in with their bracelets and read the instructions, they would see a randomly selected post from one of their Facebook or other social media profiles. They would have the option to press the button in front of them to “own up” to a quote as their own, flashing their card green so everyone could see and allow the quote to then disappear. Or, if no one claimed it, the quote was left stuck in the background of the screen for the remainder of the game. This invited users to think about the distinction between interacting publicly in-person as opposed to online, how it felt for their data to be removed from its intended context, and to consider the history of their public online identities (touching on both our Approximation of Self and Presentation of Data themes).

In contrast, Discover Your Inner Desire was an example of a individual game aiming to get participants thinking about the contractual agreements we take on when we sign up for online services (our Data Ownership and Privacy theme). Upon checking in, a user is shown a terms of service agreement that they must accept to move on. Yet upon accepting, they are led through a series of increasingly outrageous terms without the ability to get beyond them in order to make the user inspect these agreements in greater detail and confront the generally passed-over and unknown nature of terms of service agreements.

Data

A major consideration for this show was how to receive and manage data. We began by working with a lawyer to develop a simple terms of service that promised the user that no one would have access to their data prior to its use in the show and that all data would be deleted immediately after the show.

Exhibit	Description	Data Issues	N	Note
Own Up (1)	A collaborative experience where users see anonymous online quotes and choose to own up or not.	A	2-4	
Meet Your Match (2)	A compatibility engine to determine whether another person is a romantic match	A	2	
Mirror, Mirror (3)	A mirror with personalized information and messages	P	1	private
No Application Required (4)	Data is used to determine qualifications and personality traits for a job	A	2	with actor
Memory Wall (5)	A physical data piece where users map the order of major life experiences	A	1-6	non-digital
Highly Recommended	An either/or game that slowly questions users' judgment if they deviate from their data	A	1	
Discover your Inner Desire	An ever more complex series of terms of service agreements	O	1	
Wellness Booth	A system to determine happiness levels and mental health state from personal data	A	1	private
You and Your Libido	A recommender for who you're most attracted to from Facebook friends based on porn preferences	A	1	private
Interpret the Truth	Determine what facts are true or false about the world and user	R	2	
Infovision	See how cookies track users based on likes and dislikes of the news	O	1	
In the News	Get a personalized DesignCraft news article targeted to a user	P	1	

Data Issues: P: Personalization, A: Approximation of Self, O: Data Ownership and Privacy, R: Presentation of Data



Table 1. Selected exhibits to expose audience to certain personal data use issues. N denotes the number of users an exhibit is designed for.

There was also a corresponding set of bullet points presented to the user during ticketing about how their data would be used. For this pilot iteration of the production, we chose to take no anonymous statistics or behavioral metrics from user interactions with the exhibits.

Working with major data providers (e.g., Google, Twitter, Instagram) was mostly painless. However, we did have a real problem in getting Facebook to approve our app since their approval process is not designed for any kind of offline artistic experience like this. In the end, we had to film ourselves interacting with the interfaces in a rehearsal setting to get ultimate approval.

In order to keep data safe, we strongly encrypted all personal information offered by the audience. This involved generating a key pair during the ticketing process. The public key was used to encrypt all data after the mining and processing was completed. The private keys were stored in a further encrypted database that had five associated private keys, given to our four project leads and the lead set designer. At show time, two of the five keys needed to be entered in order to unlock the private key database for that night's show.

Finally, each audience member was given an RFID bracelet upon entering the performance, after showing photo ID (Figure 2.8). The bracelet's encoded ID was associated with each user's private key, allowing each user to individually choose which exhibits to check into, unencrypting the required data, or not.

Set

The set was designed to promote opportunities for different forms of interaction. Our layout had a central living/dining room where the main scenes occurred. We then had a game

room with a bar (snacks only), a den, an art studio, a bedroom, and an office. Each room had at least one exhibit in it and the layout was meant to facilitate a mixture of private conversations while also letting people see what was going on and float between scenes. Most of our walls were made of recycled crate wood slats that allowed the audience to see through into other rooms. Besides the bedroom, which was designed to allow for closed-off conversations with Amelia, the rest of the set was designed so that the main living/dining area was visible from anywhere.

Talk back

A last feature of the experience was a talk back (Figure 2.7). Following each performance we brought our entire cast and behind-the-scenes crew out in front of the audience for a discussion. We would first ask those interested to take our survey (described below) and afterwards we would open up a dialogue with our cast and crew. This would start by our production lead explaining that the point of this piece was to build engagement around these issues and that now was their time to ask or discuss anything. From here onward, we would allow the audience to get into a queue to ask questions to our team and would leave the floor open between questions to allow questions and clarifications across audience members. After 20-30 minutes we would let the actors leave, but sometimes the conversation would continue for up to an hour.

Constraints

Given our design goals and production requirements, we found ourselves having to deal with a number of constraints in our implementation of the production. As an initial case study aimed at understanding whether immersive theater is an effective method for engaging non-technologists in data

ethics, our research methods were explicitly chosen to understand our audience's sentiments around their participation, determine what kind of audience we attracted in terms of beliefs and values, and provoke reflection and discussion. We aimed to tackle doubt from academics and computer scientists that this engagement strategy may deter people due to privacy concerns, that the audience would be skewed toward those already interested in privacy, that the scope of the show was too big to be undertaken by students, or that participants would feel indifferent about data ethics after the show. This meant the engagement value of the show was taken more seriously than devising hyper-specificity in the dialogue or exhibits for research of particular technologies or systems. It also meant that we wanted our survey to be fairly short since we were already asking for 2 hours of attention at the show and did not want to limit our study to those willing to do a lengthy research program afterwards.

Further, while we were interested in collecting anonymous statistics and taking observational notes, we also did not want our audience to feel overly exposed, being observed at every juncture. In an informal pre-survey, we found that potential participants were hesitant to share their social media for research. We wanted our terms of service to have an extreme favoring of user rights to promote trust and allow people to act candidly during the performance. The research for that reason was optional, pushing us away from any monitoring of actions during the show.

These research aims along with other production goals were most constrained by timeline. Restrictions of our funding timeline and the academic calendar combined to create a very intense schedule. This meant the script had to be finalized before we developed multiple endings (something the production team all believed would serve our improvisation goal well) and we were not able to do much in terms of pre-show audience engagements.

Research Methods

In consultation with our IRB office and in light of the constraints to achieve our engagement goals, we adopted two primary forms of data collection aimed at the audience and the production team.

To determine impact on audience members, we implemented a post-show survey (Figure 2.6). We ran six performances of Quantified Self over four days, each with 40 ticketed slots, and an approximate total attendance of 240 guests. After each performance, we invited participants in the audience to stay for a 15-minute anonymous survey. The survey consisted of 20 questions aimed at obtaining information on demographics, prior experience with data and theater, general attitudes towards data ethics, and reactions to the performance. The majority of the questions had participants rank statements on a Likert scale. In total, 179 filled out the survey, although several surveys left a section or question incomplete (as reflected in the differing numbers in the results).

Utilizing a post-survey afforded us greater participation and more accurate audience experiences. First, it allowed for participation in the production by those who were wary of shar-

ing personal data. Given the hesitance and uncertainty around who would look at the personal data that was shared, we adopted a policy where no shared personal data would be used for research. Second, since research recruitment started after the performance, audience members' experiences within the production were genuine. Without the added layer of knowingly being in a research study, they were unimpaired by the sense of being watched or recorded beyond the semi-public nature of the show. Third, given our own moral grappling with data ownership, we felt it most ethical to give audience members full control of what information they would like to share with us.

In addition to our audience evaluation, we documented the design process and performed post-production interviews with the crew to look at the educational impact and reaction to being a member of the production team (Figure 2.4). We performed 15 crew interviews (5 technical team members and 10 theatrical). Interviews lasted approximately 20 minutes. We focused questions on learning (both self-reported and content questions) and impacts of the show to topic understanding and behavioral changes.

RESULTS

Following each performance, we invited participants to join in a survey. Over the 6 shows, 179 out of 240 people took the survey. Using this tool, we evaluated this project in terms of broad engagement, education, and show elements to guide future work.

Engagement

The show itself was a success as a theater piece. While we did ultimately see this as a pilot for a future scaled and improved effort, the reception clearly highlighted the potential of such a project. All six of our shows, which offered 40 slots each, were filled within the first week. As the remainder of the results will show, the audience generally left wanting more. Every night there would be a portion of the audience who would stick around, talking, until the last allowable moment.

One primary goal of this show was to engage a broader audience in conversations of ethics. In particular, we aimed to start a dialogue between technologists and non-technologists. We had two metrics in our survey to evaluate technical background: prior technical event attendance and occupation. First, we found that 71 out of 179 (40%) of participants had never attended a previous event, panel, or talk on data or technology, suggesting that this is their first attendance to an event on a technology-related subject. To complement this high number of non-technologists, we found that of the 151 respondents that specified their career or field of study, 67 (44%) were in computer science or another STEM field. This meant over half of our audience were non-technologists.

From our demographic questions, we found that there was a significant increase in female participation compared to the 26% of women involved in technical careers [2]. For gender, 89 survey participants identified as female, 85 as male, and 1 as other. Our other demographic metrics showed low ethnic and educational diversity. We found that the majority of the participants were under 30 (66%), college educated (93%),

and white (79%). This is mostly an artifact of the place where we ran this performance, which is mostly white and educated.

Educational Results

General Impact

One of our design goals was to educate all participants in this effort, including the audience and the production crew.

Regarding the audience, we report survey results relevant to their educational sentiments and perceived content accessibility. One question asked if they felt informed, confused, motivated to learn, uncertain, upset, hopeful, or concerned following Quantified Self (checking all that applied). 51% of audience members reported feeling more informed after the experience, while 16% reported feeling more confused and the remaining 33% reported neither. The number that reported neither are likely to have been from the number of people who attended who are already highly engaged in these topics.

While it supports our goals of inter-community conversation that some people came in more informed, we would hoped the percentage would be higher. However, an encouraging sign is that 62% of the audience members reported they were "motivated to learn." So, even if our show did not give everyone enough content to feel informed, it may have provided others with the impetus to go do their own research.

Another question asked audience members to compare the show to other more common modes of getting educated. Overall, we found that on average our audience somewhat agreed our show was more accessible than reading an article (3.53/5, $\sigma = 1.53$), taking a class (3.46/5, $\sigma = 0.98$), and researching online (3.52/5, $\sigma = 1.02$). We also included an unlikely alternative—reading a privacy policy. Not surprisingly, the audience most strongly agreed the show was more accessible (4.33/5, $\sigma = 0.77$).

One limitation of our finding here is that we only asked the audience members to compare to common alternatives as opposed to novel experiences they may have in the past. We are also limited to sentiments rather than, say, subject-matter expertise developed. Because we could not do a pre-survey, any content-based learning would have been impossible to assess. Given the size of the audience, it would also have been difficult to require lengthy or difficult pre/post instruments. In retrospect, adding a pre/post focus group may have been the right balance.

As reported above, the 62% who were motivated to learn following the show also signals an opportunity for continued engagement in the topics. We asked audience members about their interest in certain specific activities following the show. As a whole, the audience strongly agreed that they "want to learn more about how companies use and share personal data" (4.173/5, $\sigma = 0.65$). They further agreed that they would "want to attend a panel, talk, or event on data privacy" (3.775/5, $\sigma = 0.88$), and "want to use tools to visualize their own data" (3.98/5, $\sigma = 0.81$). This further indicates that a performance like Quantified Self may be a strong initial engagement where further programming is planned.

Engagement with Companions

Our survey had one dedicated question probing the audience sentiments of the technology exhibits. We asked for respondents to list their favorite piece, least favorite piece, and rationale for that choice. As this question was open-ended with a short response, we received less consistent answers. Out of the 179 surveyed audience members, 68 gave a complete answer and rationale for their choice. Overall, we saw that audience members preferred exhibits that elicited reflective experiences and increased social opportunities.

We did see signs of reflection around the issues of approximation of self, data ownership and privacy, and personalization. Out of the 68 survey takers who gave a rationale for their exhibit preference, 18 brought up a self-reflective moment, 12 new thoughts/knowledge, and 4 social revelations.

Self-reflection around one's self, memories, and data, was the most cited theme for positive experiences with an exhibit. "*Favorite: highly recommended. I was interested to visualize how specific data is translated into a more general image of myself*"; "*Favorite - own up. Entertaining, felt alienated to my own words when out of context.*";

Similarly, participants also felt positive about exhibits that caused them to question the nature of data privacy. "*the own-up table because it forced me to consider whether i would share publicly things that id already shared publicly.*"; "*[Discover Your Inner Desire] was the best because I wanted to play but the user agreements were crazy! It really made me question how far I am willing to share my data.*"

Even though personalization was brought up in 7 of the 68 responses, it was often negative and focused upon the accuracy of companions, and not linked to a broader ethical stand on the use of personalization. "*[Meet Your Match] was least favorite because I thought the algorithm was inaccurate*"; "*I didn't enjoy the reading room [In The News] because the content didn't seem very personalized to the participant.*" This suggests that our show could do more work in engaging the ethical question of personalization, along with the presentation of data which did not appear at all.

An exciting result was the many people most interested in the conversations and social dynamics. Own up, the exhibit with the most social interaction, was by far the most popular with 59/148 (40%) listing it as their favorite. 20% of survey takers listed social dynamics as their rationale for choosing a favorite exhibit. For example: "*Own up was my favorite - it generated live interaction with strangers*". Another person made a general statement that "*[playing] games with others could learn a little about whats important to them*".

Additional Results

Beyond participation and education, we were interested in the particular success and failure of elements of this show.

Addressing Pre-Show Concerns

One pre-show concern for this format was the use of personalized data and potential vulnerability that would create could deter attendees. Given our informal pre-show survey, people were wary of sharing data and so we evaluated attendee trust. Prior to the show, we had specifically designed a data collection policy and language clearly displayed on our website to reassure users about the use. In our survey, we evaluated this

by asking if our production or Google was trusted more with their data and why.

Out of 179 surveys, 150 answered this question with a clear choice. 63% (95/150) said our production, 25% said Google, 3% said neither, and 9% said both. A common response for those that trusted us was that we were transparent about our intentions. Out of those who trusted our production, 16 specifically mentioned the policy as engendering that trust, while many more mentioned key components (25 brought up the promise to delete data, 10 said not selling data, 7 commented upon the usage). *“You stated you would delete everything after the show so I trusted you more. Your terms and conditions were easy to read and understand.”* Those who trusted Google cited its high level of professionalism or the fact that Google already has so much information—*“it already knows everything so what’s to”* hide?”

People leaving upset was something with which we were concerned. The scripted ending of the show had our protagonist, the humanoid AI, killed (shutdown). This pessimistic ending, we worried, may have blurred the more balanced conversation throughout; however only 14% left feeling upset as a result.

Future Changes

In our survey, we asked what users would change to improve engagement in a future run of the show. Many people wanted forms of greater interaction (41/103), including more discussion with the actors, increasing data usage by the actors and exhibits, and giving the audience more agency to change the outcomes of the story. Example respondent quotes in this vein include: *“Display audience members info in front of whole audience and let that audience members reaction influence trajectory of show.”*; *“I thought the people/actors would “know” about us, not just the machines/games”*.

Others wanted more clarity on how they were supposed to interact (15%) or deeper background information (15%) coming into or during the show: *“Made it clear how much I was supposed to affect the show. For instance, Amelia’s reminder on the bookcase. I was unsure if I would unknowingly throw a wrench in the show by interacting with it.”* *“...it would have been cool if there were little clues all around giving deeper, technical, or background info to help us interact more intentionally.”*

Common themes of frustration with exhibits were wanting even more data integration and personalization or not understanding the intent of the exhibit for why an exhibit was a person’s least favorite: *“job candidacy? (didn’t understand the mechanism of pros and cons)”* **Least: Mirror** - *didn’t use much of my info”*

Cast and Crew

In addition to impacts on the audience, we conducted 15 20-minute interviews with the cast and crew (5 from the tech team and 10 from the theater team, which represented 79% of the crew) to evaluate the learning impact upon them. We found the crew members generally felt to have gained knowledge, increased their ability to discuss the topics germane to the performance, gathered insights from experiencing the audience, and appreciated the interdisciplinary interactions [43].

DISCUSSION

Where We Landed

The final production was an enriching and successful experienced both in terms of process and learned results. However, we do not believe we achieved our ideal design. Using our heuristic from above, we plotted our estimation of the final production (represented in Figure 1.2). Here you see a number of perceived asymmetries in our execution.

To begin with, we believe the show was biased toward art rather than technology. The experience was more structured by the plot and aesthetics than the technical content. Though there were technological artifacts, artistic representations overshadowed technical ones. The on-set exhibits were often more “social commentary” than showcasing or explaining the capabilities of data-driven technologies. Where we were unable to implement highly-technical systems, whether due to lack of time or team expertise, we were forced to use either simplified algorithms or entertaining content. Thus, the overall presentation of the technology more under the conceit of art rather than representing state-of-the art technical capability.

We also felt the content was more fixed than improvisational, making this more akin to a design fiction [7] or one of Odom’s semi-scripted enactments [40]. The narrative itself was not malleable in any way besides having open intervals for discussion and interaction with the characters and a few peppered in pieces of dialogue that would be generated in the hour before the performance. The exhibits too, were less personalized and more structured, making them more like design artifacts [18] than real technologies. Again, we only had time to put sophisticated techniques behind 2 of the exhibits. This means data mining for most exhibits, for instance, used keywords rather than sophisticated NLP or auto-grouping rather than a clustering analysis. Or the personalization of an exhibit was minor—replaced names or a couple messages off of your feeds.

What Worked

There is no doubt that many of the affordances discussed were achieved by this show. Beyond the fun had by all parties, we saw the artistic engagement function as an attractor for people with different backgrounds. The balance of men and women interested in and participating in the project was high. Similarly, we saw the performance house a good balance between people with STEM and non-STEM backgrounds. The curiosity around technology mixed with the entertainment of a theater piece was enough to stir up interest quickly around our campus and city.

The show was clearly consequential to many of our audience members due to the uncertainty of how their data was being used and the dynamics of the actors and exhibits that led to surprising moments. One night we had an audience member shaking and in tears at how connected he felt with our AI character and that person proceeded to talk with our crew for nearly an hour afterwards. Another night, a woman accidentally knocked a drink out of our CEO character’s hand, which actually made her worried she would be thrown out (our crew

talked with her outside and she came back in more excited than ever). You also saw many people waiting for the scripted scenes to sneak over and use the technology, likely because they were not ready for anything unexpected that may have exposed them to others. Some of the attendees thought our personal technologies were so accurate they would ask us how we knew things after the show. One person swore our Mental Health Exhibit understood he was grieving from the loss of his step-mother.

The amount of candid, interesting discussion that occurred between people was also encouraging. There was always a significant crowd of people in attendance that would ask tough questions and often stimulate other members of the audience. Not just technologists, these active players fell across the board. Sometimes privacy enthusiasts, tech hobbyists, gregarious personalities, disrupters, or lovers of role play, these people would want to be *in the world* and instigate legitimate discussion. After the show each night, many were walking away impacted and excited, as reflected in our survey results.

Perhaps the most promising aspects of this was the reporting of how many people wanted more. The results suggest that surrounding programming including panels, educational talks, and workshops would all gain better traction if paired with such a production.

What Could be Improved

There are many things we would change in a re-run of the production and lessons that others should incorporate into their own work. Even during the production process, we recognized trade-offs and compromises being made to meet our timelines and requirements. The overarching experience could have been more detailed and coherent with diegetic engagements before and after the performance. Sending a more messages or information from the characters, being able to play with your data outside of the show, and setting the background of the company and world more could have improved comprehension and participation.

An obvious takeaway for better engagement is to let audience adapt the narrative. Using a more complex interaction design, as production companies like Blast Theory coordinate, could allow for a deep exploration of one's opinions and desired outcomes with respect to relevant and important technical issues. Planning malleable character arcs and allowing the audience to affect them through their interactions with the technology and actors would take the performance to the next level. In our next run, we also want to make the technological experience more coherent and robust. We would like to see the data experience to mirror an in-show social media platform. Audience being able to see their data and discuss with one another through digital means is likely to amplify learning and engagement. This kind of interaction with live data is already trending in artistic circles, and HCI researchers could extend this trend into fascinating scenario explorations and open up possibilities to study social dynamics of networked systems.

In terms of research, had we not been on a fast-paced production timeline, we would have started the IRB process much earlier. While we do not recommend other researchers invade on the audience engagement as to limit the trust and willingness to participate, we see benefits in adding a) anonymous behavioral statistics from the exhibits and b) a small pre/post focus group to study more in depth. Both of these were missed potentials for research. For other fields, the post-survey and talk backs could be great sites for learning political views, policy positions, attitudes about technology, and gauging user experiences.

We believe this case study could have powerful potential for further research and public advocacy. Those who use design, narrative, and performance to do research should consider the choices employed in Quantified Self. Putting together the improvisational aspect of an enactment with the robust world of a design fiction, created major affordances for discussion and reflection - on both the audience and the crew. Allowing a narrative to act as a common attractor for a diverse audience to share ideas, opens up a fascinating space for interaction. Participants in these experiences appear eager for consequence and these impromptu interactions are pregnant with insights into what people want, expect, and take from technological experiences. A production like Quantified Self is an obvious candidate for ethical training within business or university settings. It is also offers a step toward moving more methodical approaches to futures research out of academia and into an accessible, mainstream format.

CONCLUSION

In this paper, we presented and a discussed a case study using immersive theater to engage, educate, and bring together technical and non-technical participants into data ethics discussions. Aiming for a balance between technical and artistic methods, and fixed and improvised modes of interaction, we developed Quantified Self: Immersive Theater and Data Experience. We found that mixing immersive theater with interactive, social-media-driven technology exhibits created opportunities for multiple forms of engagement, although many users wanted even greater interaction and personalization. Drawing from this work, we hope to see future research incorporate elements to engage non-technologists in ethical technology design and discussions.

ACKNOWLEDGMENTS

The authors would like to thank the John S. and James L. Knight Foundation for supporting the pilot run of our performance. We'd also like to acknowledge CU-Boulder's Engineering Excellence Fund, which helped support and commission student work on this project.

REFERENCES

1. Annie I. Anton, Julia B. Earp, and Jessica D. Young. 2010. How internet users' privacy concerns have evolved since 2002. *IEEE Security & Privacy* 8, 1 (2010). <http://ieeexplore.ieee.org/abstract/document/5403147/>
2. Catherine Ashcraft, Brad McLain, and Elizabeth Eger. *Women in Tech: The Facts (2015-16 Update)*. Technical

- Report. National Center for Women and Information Technology. <https://www.ncwit.org/resources/women-tech-facts-2015-16-update>
3. James Auger. 2013. Speculative design: crafting the speculation. *Digital Creativity* 24, 1 (March 2013), 11–35. DOI: <http://dx.doi.org/10.1080/14626268.2013.767276>
 4. Solon Barocas and Helen Nissenbaum. 2009. On notice: The trouble with Notice and Consent. In *Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information*. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2567409
 5. Rebecca Bates, Judy Goldsmith, Rosalyn Berne, Valerie Summet, and Nanette Veilleux. 2012. Science Fiction in Computer Science Education. In *Proceedings of the 43rd ACM Technical Symposium on Computer Science Education (SIGCSE '12)*. ACM, New York, NY, USA, 161–162. DOI: <http://dx.doi.org/10.1145/2157136.2157184>
 6. Igor Bilogrevic and Martin Ortlieb. 2016. "If You Put All The Pieces Together...": Attitudes Towards Data Combination and Sharing Across Services and Companies. ACM Press, 5215–5227. DOI: <http://dx.doi.org/10.1145/2858036.2858432>
 7. Julian Bleecker. 2009. Design Fiction: A short essay on design, science, fact and fiction. *Near Future Laboratory* 29 (2009).
 8. Mark Blythe. 2014. Research Through Design Fiction: Narrative in Real and Imaginary Abstracts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 703–712. DOI: <http://dx.doi.org/10.1145/2556288.2557098>
 9. Rainer Bohme and Stefan Kopsell. 2010. Trained to Accept?: A Field Experiment on Consent Dialogs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 2403–2406. DOI: <http://dx.doi.org/10.1145/1753326.1753689>
 10. Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. *arXiv:1607.06520 [cs, stat]* (July 2016). <http://arxiv.org/abs/1607.06520> arXiv: 1607.06520.
 11. danah boyd and Kate Crawford. 2012. CRITICAL QUESTIONS FOR BIG DATA: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication & Society* 15, 5 (June 2012), 662–679. DOI: <http://dx.doi.org/10.1080/1369118X.2012.678878>
 12. Ben Brantley. 2011. 'Sleep No More' Is a 'Macbeth' in a Hotel - Review. *The New York Times* (April 2011). <https://www.nytimes.com/2011/04/14/theater/reviews/sleep-no-more-is-a-macbeth-in-a-hotel-review.html>
 13. David Brin. 2005. The self-preventing prophecy. Or, how a dose of nightmare can help tame tomorrow's perils. *On nineteen eighty-four. Orwell and our future* (2005), 222–230.
 14. Nicholas Confessore and Danny Hakim. 2017. Data Firm Says 'Secret Sauce' Aided Trump; Many Scoff. *The New York Times* (March 2017). <https://www.nytimes.com/2017/03/06/us/politics/cambridge-analytica.html>
 15. K. Crawford. 2016. Can an Algorithm be Agonistic? Ten Scenes from Life in Calculated Publics. *Science, Technology & Human Values* 41, 1 (jan 2016), 77–92. DOI: <http://dx.doi.org/10.1177/0162243915589635>
 16. Nicholas S. Dalton, Rebecca Moreau, and Ross K. Adams. 2016. Resistance is Fertile: Design Fictions in Dystopian Worlds. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. ACM, New York, NY, USA, 365–374. DOI: <http://dx.doi.org/10.1145/2851581.2892572>
 17. Paul Dourish and Genevieve Bell. 2014. "Resistance is futile": reading science fiction alongside ubiquitous computing. *Personal and Ubiquitous Computing* 18, 4 (April 2014), 769–778. DOI: <http://dx.doi.org/10.1007/s00779-013-0678-7>
 18. Anthony Dunne and Fiona Raby. 2013. *Speculative everything: design, fiction, and social dreaming*. MIT Press.
 19. Mostafa El-Bermawy. 2016. Your Filter Bubble is Destroying Democracy. *WIRED* (Nov. 2016). <https://www.wired.com/2016/11/filter-bubble-destroying-democracy/>
 20. Chris Eلسden, David Chatting, Abigail C. Durrant, Andrew Garbett, Bettina Nissen, John Vines, and David S. Kirk. 2017. On Speculative Enactments. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 5386–5399. DOI: <http://dx.doi.org/10.1145/3025453.3025503>
 21. Chris Eلسden, Bettina Nissen, Andrew Garbett, David Chatting, David Kirk, and John Vines. 2016. Metadating: Exploring the Romance and Future of Personal Data. ACM Press, 685–698. DOI: <http://dx.doi.org/10.1145/2858036.2858173>
 22. Tristan Harris. 2016. How Technology is Hijacking Your Mind - from a Former Insider. (May 2016). <https://journal.thriveglobal.com/how-technology-hijacks-peoples-minds-from-a-magician-and-google->
 23. Luke Hutton and Tristan Henderson. 2015. "I didn't sign up for this!": Informed consent in social network research. In *Proceedings of the 9th International AAAI Conference on Web and Social Media (ICWSM)*. <https://research-repository.st-andrews.ac.uk/handle/10023/6691>

24. David Kirby. 2010. The Future is Now: Diegetic Prototypes and the Role of Popular Films in Generating Real-world Technological Development. *Social Studies of Science* 40, 1 (Feb. 2010), 41–70. DOI: <http://dx.doi.org/10.1177/0306312709338325>
25. David A Kirby. 2011. *Lab coats in Hollywood: Science, scientists, and cinema*. MIT Press.
26. Jeff Larson Lauren Julia Angwin Kirchner, Surya Mattu. 2016. Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And it's Biased Against Blacks. (May 2016). <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
27. Joshua A. Kroll, Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, and Harlan Yu. 2016. *Accountable Algorithms*. SSRN Scholarly Paper ID 2765268. Social Science Research Network, Rochester, NY. <http://papers.ssrn.com/abstract=2765268>
28. Jamiles Lartey. 2016. Predictive policing practices labeled as 'flawed' by civil rights coalition. *The Guardian* (Aug. 2016). <https://www.theguardian.com/us-news/2016/aug/31/predictive-policing-civil-rights-coalition-aclu>
29. Joseph Lindley and Paul Coulton. 2015. Back to the Future: 10 Years of Design Fiction. In *Proceedings of the 2015 British HCI Conference (British HCI '15)*. ACM, New York, NY, USA, 210–211. DOI: <http://dx.doi.org/10.1145/2783446.2783592>
30. Joseph Lindley, Dhruv Sharma, and Robert Potts. 2014. Anticipatory Ethnography: Design fiction as an input to design ethnography. In *Ethnographic Praxis in Industry Conference Proceedings*, Vol. 2014. Wiley Online Library, 237–253.
31. Conor Linehan, Ben J. Kirman, Stuart Reeves, Mark A. Blythe, Joshua G. Tanenbaum, Audrey Desjardins, and Ron Wakkary. 2014. Alternate Endings: Using Fiction to Explore Design Futures. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems (CHI EA '14)*. ACM, New York, NY, USA, 45–48. DOI: <http://dx.doi.org/10.1145/2559206.2560472>
32. Ewa Luger, Stuart Moran, and Tom Rodden. 2013. Consent for All: Revealing the Hidden Complexity of Terms and Conditions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2687–2696. DOI: <http://dx.doi.org/10.1145/2470654.2481371>
33. Jennifer Mankoff, Jennifer A. Rode, and Haakon Faste. 2013. Looking Past Yesterday's Tomorrow: Using Futures Studies Methods to Extend the Research Horizon. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 1629–1638. DOI: <http://dx.doi.org/10.1145/2470654.2466216>
34. Aaron Marcus. 2006. CHI at the Movies and on Tv. *interactions* 13, 3 (May 2006), 54–ff. DOI: <http://dx.doi.org/10.1145/1125864.1125896>
35. Aaron Marcus. 2013. The History of the Future: Sci-fi Movies and HCI. *interactions* 20, 4 (July 2013), 64–67. DOI: <http://dx.doi.org/10.1145/2486227.2486240>
36. Aaron Marcus. 2015. The Past 100 Years of the Future: HCI and User-experience Design in Science-fiction Movies and Television. In *SIGGRAPH Asia 2015 Courses (SA '15)*. ACM, New York, NY, USA, 15:1–15:26. DOI: <http://dx.doi.org/10.1145/2818143.2818151>
37. Aleecia McDonald and Lorrie Faith Cranor. 2010. *Beliefs and Behaviors: Internet Users' Understanding of Behavioral Advertising*. SSRN Scholarly Paper ID 1989092. Social Science Research Network, Rochester, NY. <http://papers.ssrn.com/abstract=1989092>
38. Omar Mubin, Mohammad Obaid, Wolmet Barendregt, Simeon Simoff, and Morten Fjeld. 2015. Science Fiction and the Reality of HCI: Inspirations, Achievements or a Mismatch. In *Proceedings of the Annual Meeting of the Australian Special Interest Group for Computer Human Interaction (OzCHI '15)*. ACM, New York, NY, USA, 670–672. DOI: <http://dx.doi.org/10.1145/2838739.2838835>
39. Jeff Nisker, Douglas K. Martin, Robyn Bluhm, and Abdallah S. Daar. 2006. Theatre as a public engagement tool for health-policy development. *Health Policy* 78, 2-3 (Oct. 2006), 258–271. DOI: <http://dx.doi.org/10.1016/j.healthpol.2005.10.009>
40. William Odom, John Zimmerman, Scott Davidoff, Jodi Forlizzi, Anind K. Dey, and Min Kyung Lee. 2012. A Fieldwork of the Future with User Enactments. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. ACM, New York, NY, USA, 338–347. DOI: <http://dx.doi.org/10.1145/2317956.2318008>
41. Paul Ohm. 2010. Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA law review* 57 (2010), 1701. http://papers.ssrn.com/sol3/Papers.cfm?abstract_id=1450006
42. Frank Pasquale. 2015. *The black box society: The secret algorithms that control money and information*. Harvard University Press.
43. Michael Skirpan, Jacqueline Cameron, and Tom Yeh. 2018. Quantified Self: An Interdisciplinary Immersive Theater Project Supporting a Collaborative Learning Environment for CS Ethics. *ACM SIGCSE Bulletin* (2018).
44. M. Skirpan and T. Yeh. 2017. Designing a Moral Compass for the Future of Computer Vision Using Speculative Analysis. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 1368–1377. DOI: <http://dx.doi.org/10.1109/CVPRW.2017.179>

45. Olivia Solon. 2017. 'This oversteps a boundary': teenagers perturbed by Facebook surveillance. *The Guardian* (May 2017).
<https://www.theguardian.com/technology/2017/may/02/facebook-surveillance-tech-ethics>
46. Bruce Sterling. 2009. Design Fiction. *Interactions* 16, 3 (May 2009), 20–24. DOI:
<http://dx.doi.org/10.1145/1516016.1516021>
47. Latanya Sweeney. 2013. Discrimination in online ad delivery. *Queue* 11.3 (2013), 10. <http://dataprivacylab.org/projects/onlineads/1071-1.pdf>
48. Joshua Tanenbaum. 2014. Design Fictional Interactions: Why HCI Should Care About Stories. *interactions* 21, 5 (Sept. 2014), 22–23. DOI:
<http://dx.doi.org/10.1145/2648414>
49. Eran Toch, Yang Wang, and Lorrie Faith Cranor. 2012. Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Modeling and User-Adapted Interaction* 22, 1-2 (April 2012), 203–220. DOI:
<http://dx.doi.org/10.1007/s11257-011-9110-z>
50. Zeynep Tufekci. 2015. Algorithmic Harms beyond Facebook and Google: Emergent Challenges of Computational Agency. *Colorado Technology Law Journal* 13.2 (2015), 203–218. <http://ctlj.colorado.edu/wp-content/uploads/2015/08/Tufekci-final.pdf>
51. John Vines, Tess Denman-Cleaver, Paul Dunphy, Peter Wright, and Patrick Olivier. 2014. Experience Design Theatre: Exploring the Role of Live Theatre in Scaffolding Design Dialogues. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 683–692. DOI:
<http://dx.doi.org/10.1145/2556288.2556960>
52. Jeffrey Warshaw, Tara Matthews, Steve Whittaker, Chris Kau, Mateo Bengualid, and Barton A. Smith. 2015. Can an Algorithm Know the "Real You"? Understanding People's Reactions to Hyper-personal Analytics Systems. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 797–806. DOI:
<http://dx.doi.org/10.1145/2702123.2702274>
53. John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 493–502.
<http://dl.acm.org/citation.cfm?id=1240704>